EC500 Enterprise

Class Series

2.5” SATA SSD

Specification

PRODUCT OVERVIEW

Dataram’s Enterprise Class 500 Series SSDs are well-balanced solid-state disk (SSD) drives with standard 2.5” form factor and great performance. Designed in SATA 6 Gb/s interface, the SSDs are able to deliver exceptional read/write speed, making them the ideal companion for mixed-use heavy-loading industrial or server operations.

The EC500 Series utilizes 3D TLC NAND for higher capacity up to 1,920GB and provides more power efficiency than 2D NAND. Regarding reliability, the EC500 is implemented with LDPC (Low Density Parity Check) ECC engine to extend SSD endurance and increase data reliability. In addition, the drives come with various implementations including powerful hardware ECC engine, power saving modes, wear leveling, flash block management, S.M.A.R.T., TRIM, and DataDefender™.

In terms of security, Advanced Encryption Standard (AES) and Trusted Computing Group (TCG) Opal ensure data security and provide users with a peace of mind knowing their data is safeguarded against unauthorized use at all times. Furthermore, with End-to-End Data Protection, data integrity can be assured at multiple points in the path to enable reliable delivery of data transfers.

FUNCTIONAL BLOCK DIAGRAM

The EC500 Series includes a single-chip SATA 6 Gb/s and the flash media. The controller integrates the flash management unit to support multi-channel, multi-bank flash arrays.

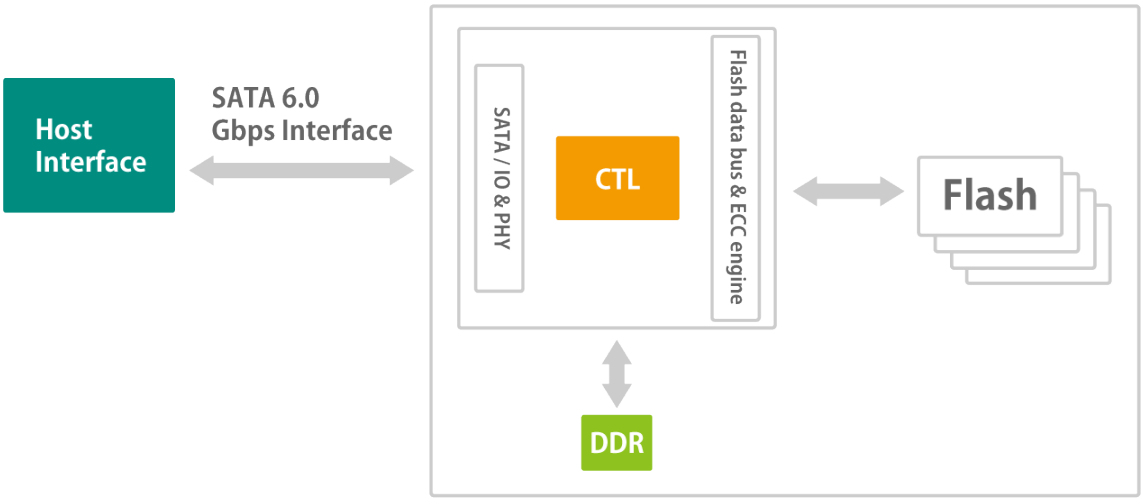

EC500S5 Enterprise Class SSD

2.5” SATA Solid State Disk

- **Compliance with SATA Revision 3.2**
  - SATA 6 Gb/s interface
  - Backward compatible with SATA 1.5 and 3 Gb/s interfaces
  - ATA command set-4 (ACS-4)
- **Capacity**
  - 240, 480, 960, 1920 GB
- **Performance***
  - Burst read/write: 600 MB/sec
  - Sequential read: Up to 560 MB/sec
  - Sequential write: Up to 505 MB/sec
  - Random read (4K): Up to 97,000 IOPS
  - Random write (4K): Up to 87,000 IOPS
- **Flash Management**
  - Low-Density Parity-Check (LDPC) Code
  - Global Wear Leveling
  - Flash bad-block management
  - Flash Translation Layer: Page Mapping
  - S.M.A.R.T.
  - DataDefender™
  - ATA Secure Erase
  - Device Sleep
  - TRIM
  - Hyper Cache Technology
  - Over-Provisioning
  - DataRAID™
  - SMART Read Refresh™
- **DRAM Cache for Enhanced Random Performance**
- **Endurance (in drive writes per day: DWPD)**
  - 240 GB: 1.42 DWPD
  - 480 GB: 1.5 DWPD
  - 960 GB: 1.37 DWPD
  - 1920 GB: 1.39 DWPD
- **Temperature Range**
  - Operating:

    Standard: 0°C to 70°C Storage: -

    40°C to 100°C
- **Supply Voltage**
  - 5.0 V ± 10%
- **Power Consumption***
  - Active mode: 425 mA
  - Idle mode: 110 mA
- **Connector Type**
  - 7-pin SATA signal connector
  - 15-pin SATA power connector
- **Form Factor**
  - 2.5”
  - Dimensions: 100.00 x 69.85 x 6.90, unit: mm
  - Net Weight: 63 g ± 5%
- **Shock & Vibration****
  - Shock: 1,500 G
  - Vibration: 15 G
- **NAND Flash Type:** 3D TLC (BiCS3)
- **MTBF:** >3,000,000 hours
- **SATA Power Management Modes**
- **Reliability**
  - Thermal Sensor
  - End-to-End Data Protection
- **Security**
  - AES 256-bit hardware encryption
  - Trusted Computing Group (TCG) Opal 2.0 (optional)
- **RoHS Compliant**
- **5 Year Conditional Warranty**

PRODUCT ORDERING PART NUMBERS

Capacity	Product Number	SED
240 GB	EC500S5AT/240G	AES256
480 GB	EC500S5AT/480G	AES256
960 GB	EC500S5AT/960G	AES256
1.92 TB	EC500S5AT/1.9T	AES256
240 GB	EC500S5PT/240G	w/TCG OPAL
480 GB	EC500S5PT/480G	w/TCG OPAL
960 GB	EC500S5PT/960G	w/TCG OPAL
1.92 TB	EC500S5PT/1.9T	w/TCG OPAL

MECHICAL SPECIFICATION

2.5” SSD physical dimensions and weight

Capacity	Height (mm)	Width (mm)	Length (mm)	Weight (gram)
All	7.00 +0.2/-0.5	69.85 ± 0.25	Max 100.45	Max 63g

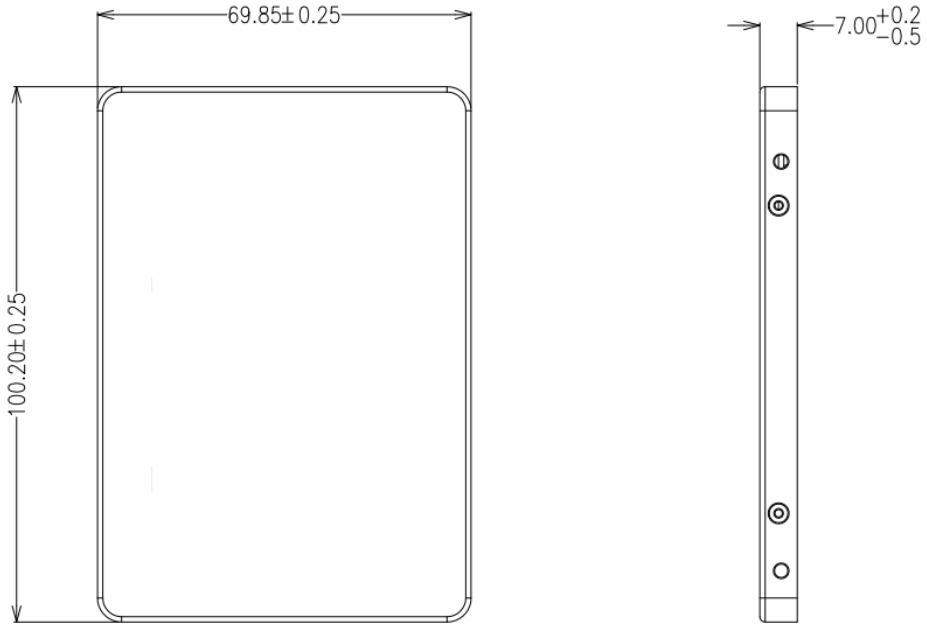

PERFORMANCE SPECIFICATION

Capacity	Data Transfer Speed (R / W)—MB/s Up to			
	Seq. Read 1 MB/s	Seq. Write 1 MB/s	Random Read 2 IOPS (4K)	Random Write 2 IOPS (4K)
240 GB	560	495	96K	79K
480 GB	560	500	97K	87K
960 TB	560	505	97K	81K
1.92 TB	560	495	97K	86K

Note:
Results may differ from various flash configurations or host system setting.
1) Sequential performance is based on CrystalDiskMark 5.2.1 with file size 1,000MB.
2) Random performance measured using IO Meter with Queue Depth 32.

ENDURANCE RATING

Capacity	TBW	DWPD
240 GB	620	1.42
480 GB	1314	1.5
960 TB	2397	1.37
1.92 TB	4793	1.39

Note:
1) This estimation complies with JEDEC JESD-219, enterprise endurance workload of random data with payload size distribution.
2) Flash vendor guaranteed 3D NAND TLC P/E cycle: 3K
3) WAF may vary from capacity, flash configurations and writing behavior on each platform.
4) 1 Terabyte = 1,024 GB
5) DWPD (Drive Write Per Day) is calculated based on the number of times that user overwrites the entire capacity of an SSD per day of its lifetime during the warranty period. (3D NAND TLC warranty: 5 years)

OPERATING VOLTAGE

Rating	Value	Unit
Supply Voltage	5V ± 10%	V

POWER CONSUMPTION

Capacity/Mode	240 GB	480 GB	960 GB	1.92 TB
Active (mA)	385	390	390	410
Idle (mA)	100	100	100	105

Note:
1) All values are typical and may vary depending on flash configurations or host system settings.
2) Active power is an average power measurement performed using CrystalDiskMark with 128KB sequential read/write transfers

ENVIROMENTAL SPECIFICATION

Environment	Specifications
Temperature	0°C to 70°C (Operating)
	-40°C to 100°C (Non-operation)
Vibration	Operation: 7.69(Grms), 20~2000(Hz)/random (compliant with MIL-STD-810G) Non-operation: 4.02(Grms), 15~2000(Hz)/random (compliant with MIL-STD-810G)
Shock	Operation: Acceleration, 50(G)/11(ms)/half sine (compliant with MIL-STD-202G) Non-operation: Acceleration,1,500(G)/0.5(ms)/half sine (compliant with MIL-STD-883K)

QUALITY AND RELIABILITY SPECIFICATION

Mean Time Between Failure (MTBF): > 3,000,000 hours

COMPLIANCE SPECIFICATION

All EC500 2.5” SSDs are compliant with the following standards and regulations:

- UL
- CE
- FCC
- RoHS
- MIL-STD-810G

PIN DESCRIPTION AND ASSIGNMENT

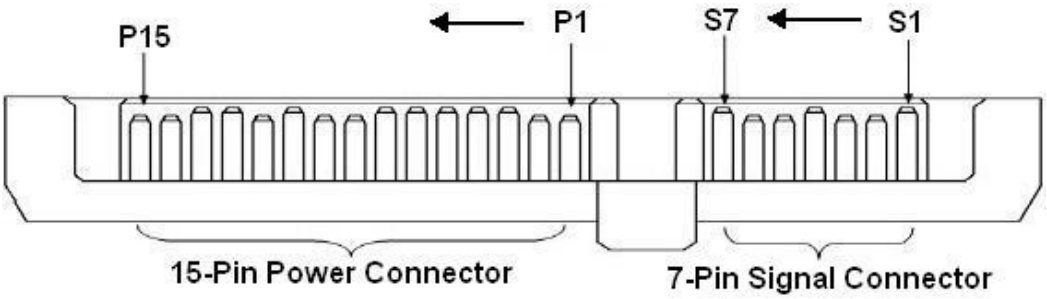

Pin	Signal Name	Description
S1	GND	2nd mate
S2	A+	Differential Signal Pair for Receiver
S3	A-	Differential Signal Pair for Receiver
S4	GND	2nd mate
S5	B-	Differential Signal Pair for Receiver
S6	B+	Differential Signal Pair for Receiver
S7	GND	2nd mate

Pin	Symbol	Description
P1	Not Used (3.3V)	N/A
P2	Not Used (3.3V)	N/A
P3	Not Used (3.3V P recharge)	
P4	GND	1st mate
P5	GND	2nd mate
P6	GND	
P7	5V P recharge	5V Power
P8	5V P recharge	5V Power
P9	5V P recharge	
P10	GND	
P11	Reserved	
P12	GND	
P13	Not Used (12V P recharge)	N/A
P14	Not Used(12V)	
P15	Not Used(12V)	

FLASH MANAGEMENT

Error Correction/Detection

The EC500 implements a hardware ECC scheme, based on the Low-Density Parity Check (LDPC). LDPC is a class of linear block error correcting code which has apparent coding gain over BCH code because LDPC code includes both hard decoding and soft decoding algorithms. With the error rate decreasing, LDPC can extend SSD endurance and increase data reliability while reading raw data inside a flash chip.

Bad Block Management

Current production technology is unable to guarantee total reliability of NAND flash memory array. When a flash memory device leaves factory, it comes with a minimal number of initial bad blocks during production or out-of-factory as there is no currently known technology that produce flash chips free of bad blocks. In addition, bad blocks may develop during program/erase cycles. Since bad blocks are inevitable, the solution is to keep them in control. The EC500 flash devices are programmed with ECC, page mapping technique and S.M.A.R.T to reduce invalidity or error. Once bad blocks are detected, data in those blocks will be transferred to free blocks and error will be corrected by designated algorithms.

Global Wear Leveling

Flash memory devices differ from Hard Disk Drives (HDDs) in terms of how blocks are utilized. For HDDs, when a change is made to stored data, like erase or update, the controller mechanism on HDDs will perform overwrites on blocks. Unlike HDDs, flash blocks cannot be overwritten and each

P/E cycle wears down the lifespan of blocks gradually. Repeatedly program/erase cycles performed on the same memory cells will eventually cause some blocks to age faster than others. This would bring flash storages to their end of service term sooner. Global wear leveling is an important mechanism that levels out the wearing of all blocks so that the wearing-down of all blocks can be almost evenly distributed. This will increase the lifespan of SSDs.

Flash Translation Layer – Page Mapping

Page mapping is an advanced flash management technology whose essence lies in the ability to gather data, distribute the data into flash pages automatically, and then schedule the data to be evenly written. Page-level mapping uses one page as the unit of mapping. The most important characteristic is that each logical page can be mapped to any physical page on the flash memory device. This mapping algorithm allows different sizes of data to be written to a block as if the data is written to a data pool and it does not need to take extra operations to process a write command. Thus, page mapping is adopted to increase random access speed and improve SSD lifespan, reduce block erase frequency, and achieve optimized performance and lifespan.

ATA Secure Erase

ATA Secure Erase is an ATA disk purging command currently embedded in most of the storage drives. Defined in ATA specifications, (ATA) Secure Erase is part of Security Feature Set that allows storage drives to erase all user data areas. The erase process usually runs on the firmware level as most of the ATA-based storage media currently in the market are built-in with this command. ATA Secure Erase can securely wipe out the user data in the drive and protects it from malicious attack.

DataDefender™

The EC500 DataDefender combines both firmware and hardware mechanisms to ensure data integrity. When power disruption occurs, the hardware mechanism will notice and trigger the controller to run multiple write-to-flash cycles to store data. Then the firmware will check that the data was correctly written to the NAND flash after the power disruption, preventing data loss.

Note: The controller unit of this product model is designed with a DRAM as a write cache for improved performance and data efficiency. Though unlikely to happen in most cases, the data cached in the volatile DRAM might be potentially affected if a sudden power loss takes place before the cached data is flushed into non-volatile NAND flash memory.

Trim

TRIM is a SATA command that helps improve the read/write performance and efficiency of solid-state drives (SSD). The command enables the host operating system to inform SSD controller which blocks contain invalid data, mostly because of the erase commands from host. The invalid will be discarded permanently and the SSD will retain more space for itself.

SATA Power Management

By complying with SATA 6 Gb/s specifications, the SSD supports the following SATA power saving modes:

- ACTIVE: PHY ready, full power, Tx & Rx operational
- PARTIAL: Reduces power, resumes in under 10 µs (microseconds)
- SLUMBER: Reduces power, resumes in under 10 ms (milliseconds)
- HIPM: Host-Initiated Power Management
- DIPM: Device-Initiated Power Management
- AUTO-SLUMBER: Automatic transition from partial to slumber.
- Device Sleep (DevSleep or DEVSLP): PHY powered down; power consumption ≦ 5 mW; host assertion time ≦ 10 ms; exit timeout from this state ≦ 20 ms (unless specified otherwise in SATA Identify Device Log).

Note: The behaviors of power management features would depend on host/device settings.

Device Sleep (DevSleep or DEVSLP) Mode

Device Sleep is a feature that allows SATA devices to enter a low power mode by designating a particular pin as DEVSLP signal with an aim to reducing power consumption.

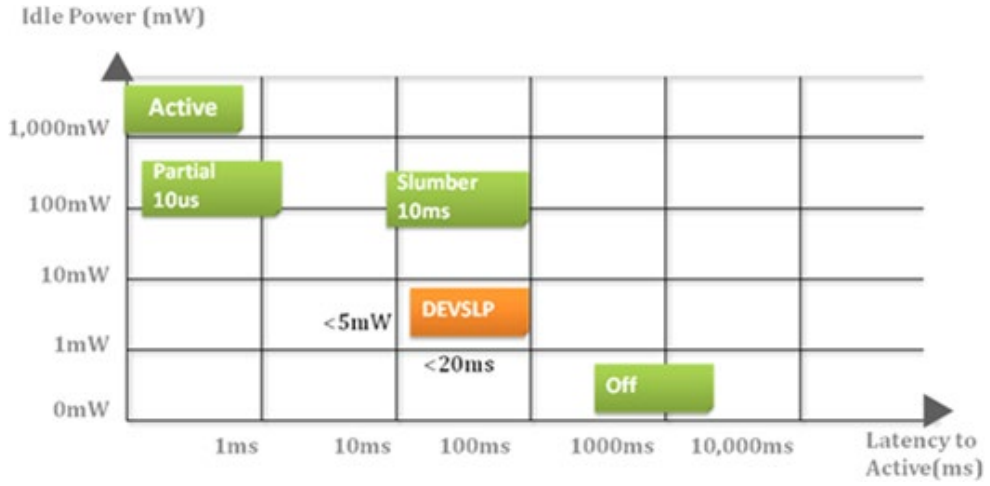

Hyper Cache Technology

The EC500 proprietary Hyper Cache technology uses a portion of the available capacity as SLC (1bit-per- cell) NAND flash memory, called Hyper cache mode. When data is written to SSD, the firmware will direct the data to Hyper Cache mode, providing excellent performance to handle various scenarios in industrial use.

Over-Provisioning

Over-Provisioning (OP) is a certain portion of the SSD capacity exclusively for increasing Garbage Collection (GC) efficiency, especially when the SSD is filled to full capacity or performs a heavy mixed-random workload. OP has the advantages of providing extended life expectancy, reliable data integrity, and high sustained write performance.

DataRAID™

The EC500 DataRAID algorithm applies an additional level of protection and error-checking. Using this algorithm, a certain amount of space is given over to aggregating and resaving the existing parity data used for error checking. So, in the event that data becomes corrupted, the parity data can be compared to the existing uncorrupted data and the content of the corrupted data can be rebuilt.

SMART Read Refresh™

The EC500 SMART Read Refresh plays a proactive role in avoiding read disturb errors from occurring to ensure health status of all blocks of NAND flash. Developed for read-intensive applications in particular, SMART Read Refresh is employed to make sure that during read operations, when the read operation threshold is reached, the data is refreshed by re-writing it to a different block for subsequent use.

SECURITY & RELIABILITY FEATURES

Advanced Encryption Standard

Advanced Encryption Standard (AES) is a specification for the encryption of electronic data. AES has been adopted by the U.S. government since 2001 to protect classified information and is now widely implemented in embedded computing applications. The AES algorithm used in software and hardware is symmetric so that encrypting/decrypting requires the same encryption key. Without the key, the encrypted data is inaccessible to ensure information security.

Notably in flash memory applications, AES 256-bit hardware encryption is the mainstream to protect sensitive or confidential data. The hardware encryption provides better performance, reliability, and security than software encryption. It uses a dedicated processor, which is built inside the controller, to process the encryption and decryption. This enormously shortens the processing time and makes it efficient.

TCG OPAL (Optional)

Developed by the Trusted Computing Group (TCG), an organization whose members work together to formulate industry standards, Opal is a set of security specifications used for applying hardware- based encryption to storage devices.

Hardware encryption has many advantages. First of all, it transfers the computational load of the encryption process to dedicated processors, reducing the stress on the host system's CPU. In addition, storage devices complying with Opal specifications are self-encryption devices. Opal specifications also feature boot authentication. When the drive is being accessed, the shadow MBR will request the drive password at boot. The drive will only unlock and decrypt if the correct password is supplied. The other feature is LBA-specific permissions. Users are assigned different permissions for LBA ranges created by the device administrator. Each LBA range is password-protected and can only be accessed by users with the correct key to perform permitted actions (read/write/erase).

Thermal Sensor

The EC500 Thermal Sensor monitors the temperature of SSD devices via S.M.A.R.T. commands. With Thermal Throttling Technology, the host system can take preventive actions by monitoring temperature changes of SSD.

End-to-End Data Protection

End-to-End Data Protection is a feature implemented in the EC500 SSD products that extends error control to cover the entire path from the host computer to the drive and back, and ensure data integrity at multiple points in the path to enable reliable delivery of data transfers. Unlike ECC which does not exhibit the ability to determine the occurrence of errors throughout the process of data transmission, End-to-End Data Protection allows SSD controller to identify an error created anywhere in the path and report the error to the host computer before it is written to the drive. This error- checking and error-reporting mechanism therefore guarantees the trustworthiness and reliability of the SSD.

SOFTWARE INTERFACE

Commands Set

This section defines the software requirements and the format of the commands the host sends to the EC500. Commands are issued to the EC500 by loading the required registers in the command block with the supplied parameters, and then writing the command code to the Command register.

Command Set

Code	Command	Code	Command
E5h	CHECK POWER MODE	F4h	SECURITY ERASE UNIT
06h	DATA SET MANAGEMENT	F5h	SECURITY FREEZE LOCK
92h	DOWNLOAD MICROCODE	F1h	SECURITY SET PASSWORD
90h	EXECUTE DEVICE DIAGNOSTIC	F2h	SECURITY UNLOCK
E7h	FLUSH CACHE	70h	SEEK
EAh	FLUSH CACHE EXT	EFh	SET FEATURES
ECh	IDENTIFY DEVICE	C6h	SET MULTIPLE MODE
E3h	IDLE	E6h	SLEEP
E1h	IDLE IMMEDIATE	B0h	SMART
91h	INITIALIZE DEVICE PARAMETERS	E2h	STANDBY
E4h	READ BUFFER	E0h	STANDBY IMMEDIATE
C8h	READ DMA	E8h	WRITE BUFFER
25h	READ DMA EXT	CAh	WRITE DMA
60h	READ FPDMA QUEUED	35h	WRITE DMA EXT
C4h	READ MULTIPLE	3Dh	WRITE DMA FUA EXT
29h	READ MULTIPLE EXT	61h	WRITE FPDMA QUEUED
2Fh	READ LOG EXT	3Fh	WRITE LOG EXT
47h	READ LOG DMA EXT	57h	WRITE LOG DMA EXT
20h	READ SECTOR	C5h	WRITE MULTIPLE
24h	READ SECTOR EXT	39h	WRITE MULTIPLE EXT
40h	READ VERIFY SECTORS	CEh	WRITE MULTIPLE FUA EXT
42h	READ VERIFY SECTORS EXT	30h	WRITE SECTOR
10h	RECALIBRATE	34h	WRITE SECTOR EXT
F6h	SECURITY DISABLE PASSWORD	45h	WRITE UNCORRECTABLE EXT
F3h	SECURITY ERASE PREPARE		

Trusted Computing Feature Set

Code	Command	Code	Command
5Ch	TRUSTED RECEIVE	5Eh	TRUSTED SEND
5Dh	TRUSTED RECEIVE DMA	5Fh	TRUSTED SEND DMA

Note: This feature set is only applicable to products implemented with AES and Opal functions.

SMART

S.M.A.R.T. is an abbreviation for Self-Monitoring, Analysis and Reporting Technology, a self- monitoring system that provides indicators of drive health as well as potential disk problems. It serves as a warning for users from unscheduled downtime by monitoring and displaying critical drive information. Ideally, this should allow taking proactive actions to prevent drive failure and make use of
S.M.A.R.T. information for future product development reference.

The EC500 devices use the standard SMART command B0h to read data out from the drive to activate our S.M.A.R.T. feature that complies with the ATA/ATAPI specifications. S.M.A.R.T. Attribute IDs shall include initial bad block count, total later bad block count, maximum erase count, average erase count, power on hours and power cycle. When the S.M.A.R.T. Utility running on the host, it analyzes and reports the disk status to the host before the device reaches in critical condition.

Note: Attribute IDs may vary from product models due to various solution design and supporting capabilities.

The EC500 memory products come with S.M.A.R.T. commands and subcommands for users to obtain information of drive status and to predict potential drive failures. Users can take advantage of the following commands/subcommands to monitor the health of the drive.

Code	SMART Subcommand
D0h	READ DATA
D1h	READ ATTRIBUTE THRESHOLDS
D2h	ENABLE/DISABLE ATTRIBUTE AUTOSAVE
D4h	EXECUTE OFF-LINE IMMEDIATE
D5h	SMART READ LOG
D6h	SMART WRITE LOG
D8h	ENABLE OPERATIONS
D9h	DISABLE OPERATIONS
DAh	RETURN STATUS

General SMART attribute structure

Byte	Description
0	ID (Hex)
1 – 2	Status Flag
3	Value
4	Worst
5*-11	Raw Data

*Byte 5: LSB

SMART attribute ID list

ID (Hex)	Attribute Name
9 (0x09)	Power-on Hours
12 (0x0C)	Power Cycle Count
163 (0xA3)	Max. Erase Count
164 (0xA4)	Avg. Erase Count
166 (0xA6)	Total Later Bad Block Count
167 (0xA7)	SSD Protect Mode (Vendor Specific)
168 (0xA8)	SATA PHY Error Count
171 (0xAB)	Program Fail Count
172 (0xAC)	Erase Fail Count
175 (0xAF)	Bad Cluster Table Count
192 (0xC0)	Unexpected Power Loss Count
194 (0xC2)	Temperature
231 (0xE7)	Lifetime Left
241 (0xF1)	Total Sectors of Write